REMARKS

The Office Action dated July 24, 2008, has been received and carefully noted. The above amendments to the claims, and the following remarks, are submitted as a full and complete response thereto.

Status of the Claims

Claims 1, 4, 7, 9, 11, 13, 16, 18-20, 22-25 and 29-34 have been amended to more particularly point out and distinctly claim the subject matter of the invention. No new matter has been added. Claims 1-9, 11-20, 22-25 and 27-34 are currently pending in the application and are respectfully submitted for consideration.

Rejection under 35 U.S.C. § 103

Claims 1-9, 11-20, 22-25 and 27-34 were rejected under 35 U.S.C. § 103(a) as being unpatentable over Bruck et al. (U.S. Patent No. 6,691,165) in view of Syvanne (U.S. Publication No. 2002/0157018) and further in view of McLaughlin et al. (U.S. Publication No. 2002/0165929). The Office Action took the position on pages 2-7 that the combination of Bruck et al., Syvanne and McLaughlin et al. teaches all of the features of the rejected claims. Applicants respectfully submit that Bruck et al., Syvanne and McLaughlin et al., both individually and in combination, fail to teach or suggest all of the features of the above-rejected claims. Reconsideration of the claims is respectfully requested.

Independent claim 1, from which claims 2-8 and 27 depend, recites a system including a network interface configured to communicate with nodes in a cluster and a

configuration subsystem operationally coupled to a remote management broker. The remote management broker is configured to distribute information between the nodes in the cluster. The system also includes a processor configured to access the cluster from a single-point, obtain information relating to at least two devices within the cluster, present the information to a user and determine network management operations to perform on the cluster. The processor is also configured to apply a configuration lock to prevent other applications from performing network management operations on the at least two devices within the cluster and to perform the determined network management operations and to determine whether the network management operations on the cluster, including said at least two devices, were applied correctly. When the network management operations were not applied correctly, the processor is configured to roll back to a successful configuration.

Independent claim 9, from which claims 11-17 and 28 depend, recites a method including accessing a cluster from a single-point, obtaining attributes relating to at least two devices within the cluster, presenting the attributes to a user and receiving input from the user relating to the attributes and determining network management operations to perform on the cluster based on the received input. The method also includes applying a configuration lock to prevent other applications from performing network management operations on the at least two devices within the cluster and performing the determined network management operations on the cluster. The method further includes determining whether the network management operations on the cluster, including said at least two

devices, were applied correctly. When the network management operations were not applied correctly, the method further includes rolling back to a successful configuration.

Independent claim 18, from which claims 19, 20, 22, 23 and 29 depend, recites a computer program embodied on a computer readable storage medium configured to control a processor to perform a process, including obtaining attributes relating to at least two devices within a cluster from a single-point, providing the attributes to a user and receiving input relating to the attributes from the user, determining network management operations to perform on the cluster based on the received input and distributing the network management operations to the at least two devices within the cluster. The process also includes applying a configuration lock to prevent other applications from performing network management operations on the at least two devices within the cluster and applying the network management operations. The process further includes determining whether the network management operations on the cluster, including said at least two devices, were applied correctly, and when the network management operations were not applied correctly, rolling back to a successful configuration.

Independent claim 24, from which claim 25 depends, recites an apparatus including obtaining means for obtaining attributes relating to at least two devices within a cluster from a single-point, providing means for providing the attributes to a user, receiving means for receiving input relating to the attributes from the user, determining means for determining network management operations to perform on the cluster based on the received input and distributing means for distributing the network management

operations to the devices within the cluster. The apparatus also includes first applying means for applying a configuration lock to prevent other applications from performing network management operations on the at least two devices within the cluster and second applying means for applying the network management operations to the devices within the cluster. The apparatus further includes determining means for determining whether the network management operations on the cluster, including said at least two devices, were applied correctly, and when the network management operations on the cluster were not applied correctly, rolling back to a successful configuration.

Independent claim 30 recites a system including network interface communicating means for communicating with nodes in a cluster, distributing means for distributing information between the nodes in the cluster, accessing means for accessing the cluster from a single-point, obtaining means for obtaining information relating to at least two devices within the cluster and presenting means for presenting the information to a user. The apparatus also includes operation determining means for determining network management operations to perform to the cluster, applying means for applying a configuration lock to prevent other applications from performing network management operations on the at least two devices within the cluster and performing means for performing the determined network management operations. The apparatus further includes correction determining means for determining whether the network management operations on the cluster, including said at least two devices, were applied correctly, and

when the network management operations were not applied correctly, rolling back to a successful configuration.

Independent claim 31, from which claims 32-34 depend, recites an apparatus including a network interface configured to communicate with nodes in a cluster and a processor. The processor is configured to access the cluster from a single-point, obtain attributes relating to at least two devices within the cluster, present the attributes to a user and receive input from the user relating to the attributes and determine network management operations to perform to the cluster. The processor is also configured to apply a configuration lock to prevent other applications from performing network management operations on the at least two devices within the cluster and to perform the determined network management operations. The apparatus is further configured to determine whether the network management operations on the cluster, including said at least two devices, were applied correctly. When the network management operations were not applied correctly, the processor is configured to roll back to a successful configuration.

As will be discussed below, Bruck et al., Syvanne and McLaughlin et al., both individually and in combination, fail to teach or suggest all of the features of the presently pending claims.

Bruck et al. generally discusses "a scalable, distributed, highly abailable, load balancing server system having multiple machines functioning as a front server layer between the network and a back-end server layer having multiple machines functioning

as Web file servers, FTP servers, or other application servers" (column 2, lines 42-47). "The operation of the servers on both layers is monitored, and when a server failure at either layer is detected, the system automatically shifts network traffic from the failed machine to one or more operational machines, reconfiguring front-layer servers as needed without interrupting operation of the server system" (column 2, lines 49-54, of Bruck et al.). "The front layer machines perform their operations without breaking network communications between clients and servers, and without rebooting of computers. In this way, the system [allegedly] provides reliable network communication in a scalable load balancing solution for server farms" (column 2, lines 63-67, of Bruck et al.).

Syvanne generally discusses "remote management of network devices" (paragraph [0001]). Syvanne seeks to "reduce the risk of loosing [sic] a management connectivity due to a misconfiguration of the network device" (see paragraph [0010]).

> [A] loss of management connectivity after a configuration changed is checked by setting up a new network connection to a device after the configuration changes are applied in the device. If it is still possible to set up a new connection and to perform the remote management after the configuration changes are applied, the new configuration is accepted for permanent use and saved permanently. If a new connection is not made and the new configuration accepted within a given time limit from starting to apply the new configuration, then the managed device automatically returns to use the previous saved configuration.

(Paragraph [0012], of Syvanne).

McLaughlin et al. generally discusses "the allocation, retention, and release of control of a programmable switch in a switched network to allow maintenance to be performed thereon" (paragraph [0005]). "[A] mechanism is provided that allows for

synchronization of online maintenance of the firmware and configuration files in the programmable switch components of a network. At the same time, the ability to direct that maintenance remains distributed among the nodes of the network" (paragraph [0012], of McLaughlin et al.).

Independent claim 1 recites, in part, determining "network management operations to perform on the cluster" and performing "the determined network management operations". Independent claims 9, 18, 24, 30 and 31, which each have their own scope, recite similar features. The Office Action took the position on page 3 that column 21, line 66, through column 22, line 13, of Bruck et al. teaches these features. Applicants respectfully disagree.

The cited section of Bruck et al. discusses that "[t]he operation of any one of the particular local monitor components 1240, 1242, 1244 can be enabled and disabled by right-clicking on the traffic signal icon for the desired component" and that "[e]nabling the monitor means that the given component (NIC, application, or ping) will be monitored" (see column 21, line 66, through column 22, line 3, of Bruck et al.). If the component is functioning properly, the signal icon is set to green, when the component has failed, the signal icon is set to red, and when the component cannot be monitored, the signal icon is set to yellow (see column 22, lines 4-9, of Bruck et al.). In other words, monitoring of components may be enabled or disabled, and the color of a signal icon may be changed according to component status.

However, nothing is cited or found in this section of Bruck et al. that teaches or suggests **determining** network management operations to perform on the cluster and **performing** the determined network management operations. Rather, it appears that Bruck et al. merely discusses enabling or disabling **component monitoring** and indicating certain **component statuses** with an icon color. As such, the cited section of Bruck et al. is completely silent as to a processor that determines which network management operations to perform and then performs said network management operations. Further, nothing is cited or found in Syvanne or McLaughlin et al. that cures these deficiencies of Bruck et al.

Independent claim 1 also recites, in part, determining "whether the network management operations on the cluster, including said at least two devices, were applied correctly, and when the network management operations were not applied correctly, the processor is configured to roll back to a successful configuration." Independent claims 9, 18, 24, 30 and 31, which each have their own scope, recite similar features. The Office Action conceded on page 3 that Bruck et al. fails to teach these features. Rather, the Office Action relied on paragraph [0012] of Syvanne to cure these deficiencies of Bruck et al. Applicants respectfully submit that Syvanne also fails to teach or suggest these features.

Per the above, the cited section of Syvanne generally discusses that a loss of management connectivity after a configuration change is checked by setting up a new network connection to a device after the configuration changes are applied in the device.

If a new connection is not made and the new configuration accepted within a given time limit from starting to apply the new configuration, then the managed device automatically returns to use the previous saved configuration. In other words, a managed device returns to a previously saved configuration when a given time period expires after a new configuration is applied.

However, Syvanne does not teach or suggest that when network management operations were not applied correctly, the **processor** rolls back to a successful configuration. Rather, it is the **managed device** in Syvanne that reverts to a previously saved configuration after a period of time. Further, Syvanne does not actually determine whether the new configuration was applied correctly. Rather, Syvanne merely relies on the managed device to revert to a previously saved configuration on its own accord after a period of time elapses. As such, the managing device does not definitively determine whether new configuration was successful. Further, nothing is cited or found in McLaughlin et al. that cures these deficiencies of Syvanne and Bruck et al.

Independent claim 1 further recites, in part, applying "a configuration lock to prevent other applications from performing network management operations on the at least two devices within the cluster". Independent claims 9, 18, 24, 30 and 31, which each have their own scope, recite similar features. The Office Action conceded on page 4 that Bruck et al. does not teach these features. Rather, the Office Action relied on paragraphs [0136]-[0144] of McLaughlin et al. to cure these deficiencies of Bruck et al.

Applicants respectfully submit that McLaughlin et al. also fails to teach or suggest these features.

The cited section of McLaughlin et al. generally discusses that "[t]he lock mechanism can be described at a high level as a coordination protocol in which nodes 12 agree to issue a 'take ownership' command to a switch 14 only after obtaining control for a lock stored in the switch 14. Under this protocol, the switch is controlled by at most one node 12 at any given time" (paragraph [0137]). The Office Action further asserted on page 4 that "[o]ne of ordinary skill in the art would have been motivated to apply the teachings of Bruck with the teachings of McLaughlin for applying a configuration lock on the at least two devices of the cluster. The motivation for doing so would have been to prevent other applications from interfering with the network management operations at the two network devices (see McLaughlin, ¶0136-¶0144).

However, the combination of Bruck et al. and McLaughlin et al. does not teach or suggest the claimed features. McLaughlin et al. seeks to prevent cooperating lock client processes to steal ownership from one another for a **single** switch while performing maintenance. This does not appear to be a **configuration lock**, as claimed, but rather a lock on a data area (see Abstract, of McLaughlin et al.). Further, while Bruck et al. discusses that the traffic assignments of multiple machines may be reconfigured, there is no teaching that these machines were intended to be locked. Applicants respectfully submit that the lack of a teaching of locking multiple devices in McLaughlin et al., combined with the lack of any teaching of locking multiple machines in Bruck et al.,

renders the combination of the cited art in an attempt to arrive at the above-recited features of the claimed invention improper. Further, nothing is cited or found in Syvanne that cures these deficiencies of McLaughlin et al. and Bruck et al.

Claims 2-4, 7, 11, 13, 16, 19, 20, 22, 23, 25, 29 and 32-34 depend from independent claims 1, 9, 18, 24 or 31 and add further features thereto. Thus, the arguments above with respect to the independent claims also apply to the dependent claims.

Per the above, Bruck et al., Syvanne and McLaughlin et al., both individually and in combination, fail to teach or suggest all of the features of the above-rejected claims under 35 U.S.C. § 103(a). Accordingly, it is respectfully submitted that the rejection is overcome and respectfully requested that the rejection be withdrawn.

Conclusion

For at least the reasons presented above, it is respectfully submitted that claims 1-9, 11-20, 22-25 and 27-34, comprising all of the currently pending claims, patentably distinguish over the cited art. Accordingly, it is respectfully requested that the claims be allowed and the application be passed to issue.

If for any reason the Examiner determines that the application is not now in condition for allowance, it is respectfully requested that the Examiner contact, by telephone, Applicants' undersigned attorney at the indicated telephone number to arrange for an interview to expedite the disposition of this application.

In the event this paper is not being timely filed, Applicants respectfully petition for an appropriate extension of time. Any fees for such an extension together with any additional fees may be charged to Counsel's Deposit Account 50-2222.

Respectfully submitted,

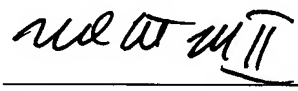

Michael A. Leonard II
Attorney for Applicants
Registration No. 60,180

Customer No. 32294
SQUIRE, SANDERS & DEMPSEY L.L.P.
14th Floor
8000 Towers Crescent Drive
Vienna, Virginia 22182-6212
Telephone: 703-720-7800
Fax: 703-720-7802

MAL:dk:sew